

Client Reference GP-270-00-US
Attorney Docket No. GOOGP024

APPLICATION FOR UNITED STATES PATENT

SYSTEMS AND METHODS FOR AUTOMATIC REPAIR AND
REPLACEMENT OF NETWORKED MACHINES

ASSIGNEE: **Google Inc.**
1600 Amphitheatre Parkway
Mountain View, CA 94043
A DELAWARE CORPORATION

ENTITY: **Large**

Jung-hua Kuo
Attorney at Law
P.O. Box 3275
Los Altos, CA 94024
Tel: (650) 988-8070
Fax:(650) 988-8090

SYSTEMS AND METHODS FOR AUTOMATIC REPAIR AND REPLACEMENT OF NETWORKED MACHINES

BACKGROUND OF THE INVENTION

1. Field of the Invention

5 [0001] The present invention relates generally to a network of computers. More specifically, systems and methods for automatic repair and replacement of computing machines are disclosed.

2. Description of Related Art

[0002] Many of today's more complex computing systems such as computer server systems often include numerous machines gathered in clusters such as a server farm. A server farm may be housed in a data center such as a collocation and may include thousands of computing machines. The machines are typically loaded with one or more server applications and allocated by use, e.g., front end, backend, etc. Over time, various machines may need to be modified, e.g., repaired, or replaced. As an example, each 15 machine may have an average service life of approximately four years. Thus, in addition to the machines that are in active use, the server farm typically also includes various free or unassigned machines that are provided as possible replacements. These replacement machines typically have various configurations such that a replacement machine with the same or similar configuration as the out of service machine being replaced may be 20 selected.

[0003] To modify or repair a machine or to replace an out of service machine with a replacement machine generally require the manual intervention by a system

administrator. The system administrator may first attempt to identify the problem, fault or exception on the machine to determine whether to modify or repair the machine. If the fault should be resolved by replacing the machine, the system administrator may select a replacement machine from a database of replacement machines to select the machine

- 5 with the same or similar configurations as the machine being replaced.

[0004] As such manual intervention is tedious, expensive, labor intensive and time consuming, it would be desirable to provide a system and method for automatically repairing or replacing a problem machine.

SUMMARY OF THE INVENTION

- 10 **[0005]** Systems and methods for automatic repair and replacement of computing machines are disclosed. It should be appreciated that the present invention can be implemented in numerous ways, including as a process, an apparatus, a system, a device, a method, or a computer readable medium such as a computer readable storage medium or a computer network wherein program instructions are sent over optical or electronic communication lines. Several inventive embodiments of the present invention are described below.
- 15

- [0006]** The system for automatic replacement of machines in a computer network may generally include a database including configuration information for the available replacement machines and a failed machine, a machine assignment module to identify 20 and assign one of the available replacement machines as a replacement machine based on a comparison of the configuration information for the failed machine to that of the available replacement machines, and a configuration module for generating configuration

data for replacement of the failed machine with the replacement machine in the computer network. The database may include configuration information for active machines in the computer network. The machine assignment module may compare certain configuration parameters such as processor speed, disk drive size, and/or amount of RAM, between the

5 failed machine and the available replacement machines. The machine assignment module may also keep track of the relative priorities of the one or more servers on each machine so as to allocate based on the relative priorities. For example, a server with higher processing power requirements may have a higher priority for more powerful machines.

10 [0007] The system may further include an installation module to cause the configuration data to take effect in at least some of the other machines in the computer network such as those that are dependent upon the failed machine. The system may also further include a detection module to detect fault in a software and/or hardware component in the machines in the computer network and a repair module to attempt to

15 repair the fault identified by the detection module in the failed machine. The system may further include a replacement module that copies data from another copy of the failed machine in the computer network, e.g., a front end server (e.g., for providing the user interface for interfacing with the end users), a load balancer (e.g., for distributing the load of a web site or other service to multiple physical servers), an index server (e.g., for indexing contents and properties of documents on the Internet or an intranet), or a cache server (e.g., for performing data replication or mirroring), into the replacement machine.

20 [0008] In another embodiment, a method for automatic replacement of machines in a computer network generally includes identifying a failed machine in the computer network, performing a lookup in a database containing configuration information for the

available replacement machines and the failed machine, identifying and assigning a replacement machine based on a comparison of the configuration information for the failed machine and the available replacement machines, and generating configuration data for replacement of the failed machine with the replacement machine in the computer network.

5 [0009] In yet another embodiment, a computer program product embodied on a computer-readable medium, the computer program product including instructions which when executed by a computer system are operable to cause the computer system to perform acts generally including identifying a failed machine in the computer network, 10 performing a lookup in a database containing configuration information for the available replacement machines and the failed machine, identifying and assigning a replacement machine based on a comparison of the configuration information for the failed machine and the available replacement machines, and generating configuration data for replacement of the failed machine with the replacement machine in the computer network.

15 [0010] These and other features and advantages of the present invention will be presented in more detail in the following detailed description and the accompanying figures which illustrate by way of example principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

20 [0011] The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements.

[0012] **FIG. 1** is a block diagram of an illustrative server farm or a clustered computer network.

[0013] **FIG. 2** is a block diagram of an illustrative automatic network machine monitor, repair and replacement system.

5 [0014] **FIG. 3** is a flowchart illustrating a process for automatically monitoring, repairing and replacing machines over a network.

[0015] **FIG. 4** illustrates an example of a computer system that can be utilized with the various embodiments of method and processing described herein.

[0016] **FIG. 5** illustrates a system block diagram of the computer system of **FIG. 4**.

10

DESCRIPTION OF SPECIFIC EMBODIMENTS

[0017] Systems and methods for automatic repair and replacement of computing machines are disclosed. The following description is presented to enable any person skilled in the art to make and use the invention. Descriptions of specific embodiments and applications are provided only as examples and various modifications will be readily apparent to those skilled in the art. The general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the invention. Thus, the present invention is to be accorded the widest scope encompassing numerous alternatives, modifications and equivalents consistent with the principles and features disclosed herein. For purpose of clarity, details relating to technical material that is known in the technical fields related to the invention have not been described in detail so as not to unnecessarily obscure the present invention.

15

20

[0018] **FIG. 1** is block diagram of an illustrative server farm or clustered computer network 100. It is noted that the computer network 100 may be located in a single location or multiple different physical locations. The computer network 100 generally includes front end servers 104 that receive and transmit user traffic 102 to the end users.

- 5 The front end servers 104 interface with backend servers 106 that provide various services such as indexing, storage, caching, searching, balancing, and/or various other processing functions, etc. It is noted that as used herein, a machine generally refers to a physical unit such as a component box containing one or more processors, disk drives, memory, power supply, etc., while a server generally refers to the one or more
- 10 applications running on the machine, i.e., a given machine may be running one or more servers. Although one exemplary server farm or clustered computer network 100 is shown, the systems and methods for automatic network machine repair and replacement as described herein may be similarly applied to various other computer network configurations.
- 15 [0019] **FIG. 2** is a block diagram of an illustrative automatic network machine monitor, repair and replacement system 120. The system may generally include a detection module 122, a repair module 124, a replacement module 126, a machine assignment module 128, a master database 130, a configuration manager 132, and a restarter module 134. Although separate components of the system 120 are shown and described for purposes of clarity, the functions of any subset or all of the various components may be integrated into a single component. The automatic network machine monitor, repair and replacement system 120 may be utilized to manage and control a computer network that includes one or more computer clusters, either locally or remotely. The system 120 facilitates in monitoring and replacing dead servers with new

replacement clones with minimal or without system administrator or other manual intervention. The system 120 may monitor various other parameters or data points and detects and generates alerts when predetermined thresholds are reached or exceeded. For example, the system 120 may monitor and detect that a disk is defective on a certain

5 machine, that there is a high average load on a given service, that there is a high collocation-wide average temperature, and the like. The detection module 122 may be configured to monitor all (or a subset of) the machines in the computer network 100. The detection module 122 may be configured with various rules to determine when software and/or hardware components of the individual machines in the computer network 100

10 may be defective, dead or otherwise malfunctioning or failed. In particular, the detection module 122 may monitor the software and/or hardware components of the individual machines in the computer network 100. For example, the detection module 122 may monitor for any defective, dead, or malfunctioning disk drives, memory, processors, etc.

[0020] If the detection module 122 detects any defective, dead, or otherwise

15 malfunctioning software or hardware components of a particular machine in the computer network 100, a token or trouble ticket identifying the particular machine and the problem is forwarded to the repair module 124. In particular, the repair module 124 receives a message from the detection module 122 that identifies which machine has a problem as well as the problem(s) detected and attempts to repair the problem(s). The repair module

20 124 may be configured with various rules to determine what repairs, if any, are to be performed. In one configuration, the repair module 124 may limit the amount or number of repairs the repair module 124 performs. For example, the repair module 124 may limit the amount of repairs for the overall network 100, for particular types of problems (e.g., memory problems, disk drive problems, etc.), for the types of machines (e.g., index

servers, cache servers, front end servers, balancers, etc.). Additionally or alternatively, the repair module 124 may be configured with rules and/or safety checks that determine whether and what repairs to perform based on the severity, history, and/or frequency of the problem, for example. In general, the detection module 122 may be configured to

- 5 report problems exhaustively while the repair module 124 may be configured to determine reasons that certain repair actions should not be performed and thus limit the triggering of repair actions.

[0021] If the detected problem is one that the repair module 124 will attempt to repair, the repair module 124 may classify the type of the problem and attempt the repair.

- 10 If the repair module 124 fails to repair the problem and/or determines that the particular machine should be replaced, e.g., without attempting to repair, the trouble ticket may be forwarded to the replacement module 126. To limit the triggering of replacement actions forwarded to the replacement module 126, the repair module 124 may be configured with safety checks. As an example, the detection module 122 may detect that a given machine
- 15 has a slightly defective disk, i.e., a defect or problem that may be fixed and does not currently require replacement, and issues a token to the repair module 124. The repair module 124 determines that the problem may be fixed and initiates a repair action. When the repair action is complete, the repair module 124 closes the repair action and forwards the token back to the detection module 122, which is continually or continuously
- 20 performing its detection loop, to determine if the problem was actually fully fixed. If the detection module 122 still detects the problem, the detection module 122 issues the token again and forwards the token to the repair module 124 again to trigger the fix action. The repair module 124 (and/or optionally the detection module 122) may keep track of the number of times and when the repair(s) were attempted and may be configured to manage

the frequency of the repair, i.e., how soon after the previous repair should the repair module 124 attempt the repair again. The repair module 124 may also be configured to limit the number of repairs of the same problem and/or limit the number of the same repairs for the particular machine, etc. When the maximum number of repair attempts is

- 5 reached, e.g., 5 repair attempts, the repair module 124 may terminate the repair cycle, for example, by modifying the token to indicate that the problem is not repairable and forwarding the modified token back to the detection module 122. The modification of the token thus converts the token to a different problem type for which the prescribed action is replacement of the machine. The machine may then be declared unusable and a
- 10 replacement action may be triggered.

[0022] To trigger a replacement action, the replacement module 126 forwards the identification of the machine to the machine assignment module 128 so as to identify a best match replacement machine from the pool of available replacement machines for the particular failed machine. The machine assignment module 128 in turn interfaces with

- 15 the master database 130 which contains a list of all the machines in the computer network, including those in the pool of available replacement machines. For each machine, the list in the master database 130 may include each machine's identification and the particular configuration, e.g., processor speed, disk drive size, and/or amount of RAM, etc. The machine assignment module 128 performs a lookup in the master
- 20 database 130 of the failed machine to be replaced to determine the configurations of the failed machine. The machine assignment module 128 then determines the best match replacement machine from the list of available replacement machines in the master database 130. Typically, each machine in the pool of available replacement machines is up and running but not yet assigned and thus available for assignment. The machine

assignment module 128 then assigns the selected best match machine as the replacement machine and executes an automatic script to copy, mirror or clone the new replacement machine with data from another copy of the machine being replaced, e.g., another copy of the same index server or same balancer, etc.

- 5 [0023] Once the new replacement machine is mirrored with data from another copy of the machine being replaced, the new replacement machine is ready to be installed in the computer network. Specifically, the configuration manager 132 generates a new configuration file (or multiple configuration files, depending on the particular implementation, for example) that takes into account replacement of the failed machine
- 10 10 with the new replacement machine in the computer network. The configuration manager 132 forwards the new configuration file to the installation or restarter module 134 and optionally also forwards a request to replace the failed machine with the new replacement machine as a redundant safety measure.

[0024] The restarter module 134 performs the installation of the replacement machine

- 15 15 in the computer network by updating or installing the configuration files of the other machines in the computer network that are in the failed machine's chain of dependency.

[0025] The restarter module 134 performs the installation of the replacement machine in the computer network by restarting the binary on the target machine so its command line reflects the newly added machine(s).

- 20 [0026] configuration files generated by the replacement module 126 are moved from Perforce (at corp) to the babysitters (in production) where the restarter modules 134 are (there is one babysitter/restarter per coloc while there is a single, central, configuration manager 132

[0027] Specifically, the failed machine's chain of dependency includes the other machines in the computer network that are affected by the replacement of the failed machine. For example, the restarter module 134 may parse the configuration files, build the hierarchy, and enforce, for example, that machines performing function A

- 5 communicate with machines performing function B, machines performing function B communicate with machines performing function C, etc. The restarter module 134 may run an automatic script that instructs the machines dependent on the failed machine to modify or update its configuration file with the new configuration file as generated by the configuration manager 132. The restarter module 134 may also cause the modified
- 10 dependent machines to restart in order for the new configuration data to take effect.

Generally, the restarter module 134 may also function as a babysitter of the machines of the computer network, e.g., by monitoring for dead or non-executing machines in general and, upon finding a dead machine, executing an automated script to restart the dead machine.

- 15 [0028] FIG. 3 is a flowchart illustrating a process 150 for automatically monitoring, repairing and replacing machines over a network. The automatic process 150 may be utilized to manage and control, either locally or remotely, a computer network that includes one or more computer clusters that may be located in one or multiple physical locations. The process 150 facilitates in the automatic monitoring and replacing of dead
- 20 servers with new replacement cloned machines with minimal or without system administrator or other manual intervention.

[0029] At block 152, all or a subset of the machines in the computer network are monitored. Various rules may be employed to facilitate in determining when software and/or hardware components of the individual machines in the computer network may be

defective, dead, failed, or otherwise malfunctioning. For example, individual components such as disk drives, memory, processors, etc. of each machine may be monitored for any defect or malfunction.

[0030] If no defect or malfunction is detected as determined at decision block 154,

5 the process returns to the monitoring at block 152. Alternatively, if a defect or malfunction is detected as determined at decision block 154, a token or trouble ticket is generated to identify the failed machine and the associated problem at block 156. At decision block 158, a determination is made as to whether to attempt to repair the problem. For example, the automatic process 150 may be configured with various rules

10 to determine what repairs, if any, are to be performed on the failed machine or the particular software or hardware component of the identified machine. For example, the amount or number of repairs may be limited for the overall computer network, for particular types of problems such as memory problems, disk drive problems, etc., and/or for the types of machines such as index servers, cache servers, front end servers,

15 balancers, etc. Additionally or alternatively, the process 150 may be configured with rules that determine whether and what repairs to perform based on the severity, history, and/or frequency of the problem, for example.

[0031] If the detected problem is one that the process 150 will attempt to repair or replace as determined at decision block 158, the type of problem (e.g., repair action or

20 replacement action) may be classified. If the prescribed action is a repair action as determined at decision block 160, the repair action is performed at block 162. The process 150 then returns to the monitoring at block 152 which determines whether the repair action performed at block 162 was successful. If the attempted repair is successful, the monitoring at block 152 would not detect the problem again. However, if the

attempted repair is unsuccessful, then the monitoring at block 152 again generates the token or trouble ticket at block 156. Note that as discussed above, the number of times and/or when previous repair action(s) were made may be monitored in order to manage the frequency of the repair (i.e., how soon after the previous repair should the repair

- 5 action be performed again) and/or whether to issue a replacement action such as when the maximum number of the repairs (e.g., repair of the same problem and/or the same repairs for the particular machine, etc.) is reached. For example, the process 150 may determine at block 158 that the maximum number, e.g., 5, of repair actions is reached, and modifies the token to indicate that the problem is not repairable (and thus may be replaced) and
- 10 forwards the token back to the monitoring at block 152.

[0032] Alternatively, if the prescribed action is replacement as determined at decision block 160, the best match replacement machine for the failed machine is identified and assigned from a pool of available replacement machines at block 166. In particular, at block 166, the configuration of the failed machine may be determined by performing a

- 15 lookup based on the identification of the failed machine in a database of all machines in the computer network and their configurations. The best match replacement machine may then be selected from a pool of available replacement machines in the database based on the configuration of the failed machine. The best match replacement machine may be the machine most similar in configuration to that of the failed machine. As an
- 20 example, the best match replacement machine may be selected to have the closest but at least the same or better configuration for each parameter, e.g., processor speed, disk drive size, and/or amount of RAM, etc. Alternatively, the best match replacement machine may be selected to have the closest but at least the same or better configuration for only certain predetermined parameters, e.g., same or better processor speed but disk drive size

may be irrelevant or may be any size greater than 20 GB despite that the failed machine has a disk drive size of 40 GB.

[0033] Next, at block 168, data from another copy of the failed machine, for example, another copy of the same index server or same balancer, is copied onto the newly identified and assigned replacement machine. As is evident, the copy of the failed machine should be operational, i.e., no fault detected. At block 170, a new configuration file is generated that takes into account the replacement of the failed machine with the new replacement machine in the computer network. At block 172, the machines in the failed machine's chain of dependency are modified, updated or installed with the new configuration file. As noted above, the failed machine's chain of dependency includes the other machines in the computer network that are affected by the replacement of the failed machine, typically those hierarchically higher in the computer network than the failed machine. Finally, at block 174, the modified dependent machines may be restarted in order for the new configuration to take effect. Depending on the configuration of the computer network and/or the replacement and its dependent machines, for example, the restart of the modified dependent machines may be staggered relative to one another so as to ensure that the computer network system remains functional during the restart, i.e., during the installation of the new replacement machine. Once the new replacement machine is installed in the computer network, automatic replacement of the failed machine is complete.

[0034] FIGS. 4 and 5 illustrate a schematic and a block diagram, respectively, of an exemplary general purpose computer system 1001 suitable for executing software programs that implement the methods and processes described herein. The architecture and configuration of the computer system 1001 shown and described herein are merely

illustrative and other computer system architectures and configurations may also be utilized.

[0035] The exemplary computer system 1001 includes a display 1003, a screen 1005, a cabinet 1007, a keyboard 1009, and a mouse 1011. The cabinet 1007 typically houses 5 one or more drives to read a computer readable storage medium 1015, a system memory 1053, and a hard drive 1055 which can be utilized to store and/or retrieve software programs incorporating computer codes that implement the methods and processes described herein and/or data for use with the software programs, for example. A CD and a floppy disk 1015 are shown as exemplary computer readable storage media readable by 10 a corresponding floppy disk or CD-ROM or CD-RW drive 1013. Computer readable medium typically refers to any data storage device that can store data readable by a computer system. Examples of computer readable storage media include magnetic media such as hard disks, floppy disks, and magnetic tape, optical media such as CD-ROM disks, magneto-optical media such as floptical disks, and specially configured hardware 15 devices such as application-specific integrated circuits (ASICs), programmable logic devices (PLDs), and ROM and RAM devices.

[0036] Further, computer readable storage medium may also encompass data signals embodied in a carrier wave such as the data signals embodied in a carrier wave carried in a network. Such a network may be an intranet within a corporate or other environment, 20 the Internet, or any network of a plurality of coupled computers such that the computer readable code may be stored and executed in a distributed fashion.

[0037] The computer system 1001 comprises various subsystems such as a microprocessor 1051 (also referred to as a CPU or central processing unit), system memory 1053, fixed storage 1055 (such as a hard drive), removable storage 1057 (such as

a CD-ROM drive), display adapter 1059, sound card 1061, transducers 1063 (such as speakers and microphones), network interface 1065, and/or printer/fax/scanner interface 1067. The computer system 1001 also includes a system bus 1069. However, the specific buses shown are merely illustrative of any interconnection scheme serving to link 5 the various subsystems. For example, a local bus can be utilized to connect the central processor to the system memory and display adapter.

[0038] Methods and processes described herein may be executed solely upon CPU 1051 and/or may be performed across a network such as the Internet, intranet networks, or LANs (local area networks) in conjunction with a remote CPU that shares a portion of 10 the processing.

[0039] While the exemplary embodiments of the present invention are described and illustrated herein, it will be appreciated that they are merely illustrative and that modifications can be made to these embodiments without departing from the spirit and scope of the invention. Thus, the scope of the invention is intended to be defined only in 15 terms of the following claims as may be amended, with each claim being expressly incorporated into this Description of Specific Embodiments as an embodiment of the invention.